The End of Behavioral Insurance

by

Glenn W. Harrison [†]

December 2023

Forthcoming, Georges Dionne (ed.) *Handbook of Insurance*
(New York: Springer, 2024, Third Edition)

ABSTRACT

Our descriptive understanding of observed insurance behavior has been enhanced by considering alternative modeling approaches, and promises to do the same to our normative evaluation of that behavior. Those alternatives come from the field of behavioral economics, which offers explicit, alternative characterizations of the way in which decisions have been made. The value of these alternatives is clear in a wide range of topics in empirical insurance, to the point where there is now no reason to debate why we need to consider them. Indeed, recognition that behavioral insurance has come to stay in our scholarship allows us to signal the end of the need to even make the case for behavioral insurance. However, that recognition does not mean that every claim from behavioral insurance is to be accepted at face value, and many are dubious or loosely applied. Much remains to be done more carefully, and much simply remains to be done.

KEYWORDS: behavioral insurance, behavioral economics, risk preferences, intertemporal risk aversion, subjective beliefs

JEL CODES: D01, D9, D81, D83

[†] Maurice R. Greenberg School of Risk Science and Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, USA. Harrison is also affiliated with the School of Economics, University of Cape Town. Comments from three reviewers and the editor are appreciated. E-mail contact: gharrison@gsu.edu.

1. Introduction

Decisions to purchase insurance should be a perfect place to see economic theory at work in general, and behavioral economics at work in particular. We have well-developed descriptive theories of the demand for insurance products, and many of these theories extend relatively easily to the insights of behavioral economics. When we turn to normative issues, however, things are not nearly so settled. The informational requirements needed to undertake welfare evaluations are already severe, and encounter subtleties when we turn to *behavioral* welfare economics.

From a theoretical perspective, one can quickly identify several "behavioral moving parts" in determining the demand for canonical insurance contracts. One set of those moving parts has to do with risk preferences, time preferences, and subjective beliefs of the (potentially) insured, another set has to do with asymmetric information between the insured and the insurance company.

Consider the first set of behavioral moving parts. The first set is *atemporal* risk aversion, which can derive from various psychological pathways, such as aversion to outcome variability and probability weighting. The second is subjective beliefs about loss probabilities, as well about non-performance risk and other basis risks when applicable. The third concerns time preferences, thinking of insurance as an explicitly time-dated contract. In many product lines, the insurance contract specifies a known premium payment *now* in the expectation that if something happens to the policyholder *over the coming year* the insurer will honor that contract and help mitigate the loss. In some important product lines, such as health insurance, the premia are typically paid in equal amounts over time. And the fourth involves the interaction of risk and time preferences, in the form of *intertemporal* risk aversion; more on this below, since it seems to be less familiar to many insurance economists.[1]

Now consider the second set of behavioral moving parts, deriving from asymmetric information. The first part is *adverse selection* on risk type, defined in terms of the perceived loss probability of the

[1] Harrison and Ng (2019) review experimental research in behavioral insurance.

insured. Here we need to be clear about whose perceptions are relevant for different decisions, and in general we want to allow there to be perceptions by the insured as well as perceptions by the insurance company (underwriters). The canonical case involves the insured having perfect knowledge of the loss probability, and the insurance company having none. Adverse selection is expected to be a particularly serious problem for health insurance, and is often mitigated by government requirements for insurance (e.g., vehicle or homeowners insurance). The second is *moral hazard*, defined as the insured choosing to engage in effort to minimize the probability and/or size of a claim. The canonical case involves the insured have costless or low-cost ways to minimize expected claim amounts, and the insurance company having no information on whether the insured made that effort. Both adverse selection and moral hazard derive from there being some unobservable trait of the insured, either their risk type or their effort level.

Section 2 clarifies what "behavioral economics," and hence behavioral insurance, refers to. Section 3 briefly reviews key concepts in risk preferences and subjective beliefs, central to any descriptive or normative evaluation of insurance behavior. Section 4 reviews a selective series of substantive issues and debates in the behavioral insurance literature. General lessons are drawn in section 5.

2. Spoiler Alert: The End of Behavioral Insurance

Behavioral economics in general suffers from being defined by many in terms of the issues that gave it life, rather than what it has become over time. The same is true of behavioral insurance. Some modern history of thought might therefore be useful, to correctly set the stage.

One of the least useful definitions of behavioral economics is that it deals with the irrationalities that characterize behavior we actually observe, in contrast to traditional economics that assumes omniscient, all-calculating agents that always follow some putative model of rational action. One problem with this definition is that it confuses commonly used models of economic behavior with rationality in general.[2] The other problem is that it confuses the reasonable use by economists of context-specific notions of ecological rationality with general, all-purpose rationality, a notion that economists arguably have no need of.[3] What this narrow definition does reflect is the valuable role that controlled laboratory experiments, and the apparent anomalies they generated, played in the founding of behavioral economics.

In the empirical insurance literature we find references to consumers behaving rationally when making choices to purchase or not, particularly when the focus is on some issue to do with adverse selection. Examples include Chiappori et al. (2006; §2), Einav et al. (2010a; p.885) or Fang and Wu (2018; p. 762). In most cases "rational" just means the use of Expected Utility Theory (EUT) or direct revealed preference. This usage is innocent, and just a technical shorthand, but we resist it for good cause.

Another less-useful definition of behavioral economics is a merging of economics and psychology. At best this definition over-simplifies, glossing the best and worst of both literatures. At worst it encourages economists to accept on face value claims from psychology that have not been demonstrated with the qualities we have come to expect in economics, such as persistence when agents face salient incentives rather than abstract hypothetical survey questions, operational definitions of core concepts, and connection to wider theories of behavior. To be sure, some of the claims from psychology do survive such demonstrations, as illustrated brilliantly by the motivating experimental design of Grether and Plott (1979). But we must not be in a position of failing to "kick the methodological tires" when being asked to incorporate some claim from psychology.

Richter et al. (2014) adopt a mix of definitions of behavioral insurance. They start by viewing it as just "adding cognitive factors" to EUT models:

> Risk-taking decisions can be highly complex and highly dependent on the specific situation of each decision maker. [...] Although EUT has some predictive power, it also seems to have many contradictions to predicted outcomes. Behavioral models attempt to add various cognitive factors into the process. Some of these might be simplifications; some

[2] For example, anyone that uses Subjective Expected Utility, and claims that it is supposed to apply in general, has simply not read Savage (1972) and his notions of "small worlds" and "large worlds" that define the domain of applicability of the formal theory.

[3] See Harrison and Ross (2023; §3.C and §4) for further discussion.

> might be awareness of a social context; and some might be recognition of one's own emotions, hopes and fears.

And after reviewing substantive contributions to the field, they take a pragmatic, agnostic view of what behavioral insurance is: "Behavioral models of insurance come in many packages. Essentially, they can act as either complements to or substitutes for classical models based on EUT." Although limited to atemporal risk preferences, this is about right as a summary of what the literature does.

I would offer an amended perspective of what behavioral insurance has now become, consistent with this pragmatic, agnostic view, but broadening it. The value of modeling alternative "behavioral moving parts" is now widely accepted. In some challenging areas, the subtleties of modeling make it distracting to take on too many possible variants, but that is just a reflection of theoretical and empirical modelers focusing on what they see as essentials as they explore certain issues. We will see many "speculative" examples of this strategy in section 4. But when we accept the value of modeling alternatives in general, we do not have to agree on which ones are relevant, or that any of them are, for specific applications. The idea that *everyone* violates EUT, Exponential Discounting or Bayes Rule is just silly and wrong when one does the empirical analyses rigorously at the individual level. Enough individuals do violate one or more of these to make it valuable to account for such behavior. Nor should one endow "the representative agent" with too much scholarly authority, since that quickly leads to confusing the average with the typical. So one does not have to kneel down before the extravagant claims of some behavioral economists about how individuals "actually behave" to accept the value of modeling variants. In that sense, the term "behavioral" should just melt away in general.

The same idea was well proposed by Camerer (2003; p.465) with respect to behavioral game theory:

> This book describes a large, and rapidly growing, body of experimental data designed to address two major criticisms of game theory: first, that game theory assumes more calculation, foresight, perceived rationality of others, and (in empirical applications) self-interest than most people are naturally capable of; and, second, that in most applied domains there is too much theorizing about how rational people would interact strategically, relative to the modest amount of empirical evidence on how they do interact.

> (No science – especially the "hard" sciences economists envy most, such as physics, chemistry and biology – has flourished without a very large dose of data-constraining theorizing.)
>
> Both criticisms can be addressed by observing how people behave in experiments in which their information and incentives are carefully controlled. These experiments test how accurately game-theoretic principles predict the behavior. When principles are not accurate, the results of the experiment usually suggest alternative principles. This dialogue between theory and observation creates an approach called "behavioral game theory," which is a formal modification of rational game theory aided by experimental evidence and psychological intuition. [...] The eventual goal is for game theorists to accept behavioral game theory as useful and necessary. When that time comes, the central ideas in this book will be part of every standard game theory book and the term "behavioral" can be shed.

Hence this is all about the end of behavioral insurance in this sense.

3. The Moving Parts of Behavioral Insurance

Insurance contracts focus attention on some of the core concepts of risk preferences and subjective beliefs in economics. In turn, behavioral economics has expanded our formal understanding of these concepts, and presumably their use when we turn to the behavioral insurance literature.

The empirical literature in behavioral insurance can be classified into three broad categories. One reason for doing this is to point out the advantages and disadvantages of each approach, and to suggest ways that hybrid approaches might mitigate these disadvantages.

The first approach is a "tops down" methodology that starts with some observed field data that has certain essential features of an experimental design, and asks what identifying restrictions are needed to make certain inferences about behavior. It does not matter if the experimental design was not the intended to aid these inferences: it might just be as simple as the customer being offered a menu of alternative contracts. Indeed, it is often just a menu of insurance contracts with the only objective differences being the deductible and the premium of each contract. The advantage of this approach, of course, is that it directly places the researcher and her inferences in the field, in the domain of naturally occurring behavior. The disadvantage is that the identifying restrictions, in terms of risk preferences and subjective beliefs of the customer, often need to be very severe indeed. One of the concerns with this literature is that it often leads to claims that some non-standard behavioral pattern or "friction" has to be at work in order to explain the observed data adequately. The risk here is that the severe identifying assumptions with respect to risk preferences and subjective beliefs might also have explained some or all of those observed data patterns. We highlight the severity of these restrictions in section 4, and link them to alternative hypotheses that could account for the observed data.

The second approach is a "bottoms up" methodology that starts with some structural theory about how insurance decisions are made, then designs experiments to allow one to identify the "behavioral moving parts" of that structural theory. There is no need for the structural theory to be limited to familiar, standard models of risk preferences or subjective beliefs, but they are often a natural starting place. The strength of this approach is that it directly connects the researcher and her inferences to a structural theory, so that there should be no ambiguity over what the resulting inferences about behavior mean. One limitation of the *application* of this approach is that it is often applied only to convenience sample of university students, even though there have long been "artefactual field experiments" doing exactly the same thing with inconvenient samples that are representative of populations (see Harrison and List (2004)). The use of "auxiliary" artefactual tasks to statistically condition inferences about behavior is the methodological contribution coming from having structural models of behavior, whether that methodological insight is then applied in the laboratory or the field.[4]

Where the "tops down" and "bottoms up" approaches sharply conflict is when we move from descriptive analyses of behavior to normative analyses of behavior (Harrison (2019)). If something is modeled as a "friction" or a "mistake," rather than a preference or a belief, we face different challenges when normatively evaluating behavior. In the former case there is a presumption that removing or overcoming the "friction" or "mistake" will improve welfare for individuals deciding whether to purchase

[4] Another limitation of the application of this approach is that researchers often use proxy measurements of the risk preferences or subjective beliefs needed, such as hypothetical surveys. In general, these are known from decades of research to be unreliable. This is a limitation of the mis-application of empirical methods, akin to the universal use of Ordinary Least Squares estimators in some fields.

insurance or not. In the later case we need to investigate further if the preference or belief provides a basis for normative inference, as stressed by Harrison and Ross (2018). But it is often the case that the preference or belief are normatively attractive, in the "consumer sovereignty" spirit of welfarism (that welfare judgements should be made on the basis of the preferences and beliefs of the affected individuals). So we quickly end up with sharply different normative implications of the "tops down" and "bottoms up" approaches.

A third approach can be thought of as a hybrid mix of the first two approaches. In this case we augment the field observations of the "tops down" approach with *priors* about preferences and beliefs from other sources. This is just recognizing "nuisance parameters" from the point of view of statistical identification, and then conditioning on them with non-degenerate priors. For example, one could simply run artefactual field experiments to estimate the preferences and beliefs of samples from the sample population. Or one could use experiments from comparable subjects, with the recognition that these are only comparable subjects, not subjects drawn from the same (target) population. The latter step alerts us to the fact that these are *priors* that the researcher has over the preferences and beliefs of the target population. We can then reasonably discuss what might make better or worse priors for these descriptive or normative inferences, but at least we are focusing on the right idea of a prior rather than magically being able to estimate the "true" preferences and beliefs of the target population, as stressed by by Harrison and Ross (2023).

A. Atemporal Risk Aversion

For atemporal risk aversion, different theories agree on what defines the risk premium, but then decompose it differently. EUT attributes all of the risk premium to aversion to variability of outcomes, measured by the non-constant marginal utility of outcomes as the level of the outcomes vary. If the risk premium is positive, and there is risk aversion, this is diminishing marginal utility and a concave utility function. Rank-Dependent Utility (RDU), due to Quiggin (1982), adds to this account of the risk premium some allowance for various forms of probability weighting, leading to decision weights on the utilities of outcomes that can differ systematically from observed or subjective probabilities. These are just two of the most important structural models, and the ones we primarily consider.

Formally, the two axioms of EUT that are of greatest concern are the Compound Independence Axiom (CIA) and the Reduction of Compound Lotteries (ROCL) axiom.[5] Many classical experimental tests of EUT were designed to test the combination of the CIA and ROCL, known now as the Mixture Independence Axiom (MIA).[6] We will see soon why it has become important to tease the CIA and ROCL apart in modern behavioral insurance: there are many applications where one wants to relax one or the other, but not both.

Using these axioms, we can formally characterize RDU in relation to EUT as relaxing the CIA while maintaining ROCL. Specifically, RDU replaces the CIA with a Comonotonic Independence axiom defined over rank-ordered outcomes, and EUT is nested within RDU. Machina (1987) and Starmer (2000) for accurate, expert statements of the early evaluation of axioms for choice over risky alternatives.

An important class of risk preferences relaxes ROCL but maintains the CIA. This is the Recursive RDU model due to Segal (1990)(1992). The basic idea is to assume that the second-stage lotteries of any

[5] Let X, Y and Z denote simple lotteries, A and B denote compound lotteries, $\succ$ express strict preference, and ~ express indifference. The CIA says that if A is the compound lottery giving the simple lottery X with probability α and the simple lottery Z with probability $(1\text{-}\alpha)$, and B is the compound lottery giving the simple lottery Y with probability α and the simple lottery Z with probability $(1\text{-}\alpha)$, then $A \succ B$ iff $X \succ Y \; \forall \alpha \in (0,1)$. So the construction of the two compound lotteries A and B has the familiar "independence axiom" cadence of the common prize Z with a common probability $(1\text{-}\alpha)$, but the implication of the CIA is only that the *ordering* of the compound and constituent simple lotteries are the same. The ROCL axiom says that A ~ X if the probabilities and prizes in X are the actuarially-equivalent probabilities and prizes from A. Thus if A is the compound lottery that pays "double or nothing" from the outcome of the lottery that pays \$10 if a coin flip is a head and \$2 if the coin flip is a tail, then X would be the lottery that pays \$20 with probability $\frac{1}{2}\times\frac{1}{2} = \frac{1}{4}$, \$4 with probability $\frac{1}{2}\times\frac{1}{2} = \frac{1}{4}$, and nothing with probability ½. From an observational perspective, one would have to see choices between compound lotteries and the actuarially-equivalent simple lottery to test ROCL: see Harrison et al. (2015a).

[6] The MIA says that the preference ordering of two simple lotteries must be the same as the two actuarially-equivalent simple lotteries derived from the two compound lotteries formed by combining a common outcome with one of the original simple lotteries, where the common outcome has the same (compound lottery) probability. That is, $X \succ Y$ iff the actuarially-equivalent simple lottery of $\alpha X + (1\text{-}\alpha)Z$ is strictly preferred to the actuarially-equivalent simple lottery of $\alpha Y + (1\text{-}\alpha)Z$, $\forall \alpha \in (0,1]$. So stated, it is clear that the MIA strengthens the CIA by making a definite statement that the constructed compound lotteries are to be evaluated in a way that is ROCL-consistent. Construction of the compound lottery in the MIA is implicit: the axiom only makes observable statements about two pairs of simple lotteries.

compound lottery are replaced by their certainty-equivalent, "throwing away" information about the second-stage probabilities before one examines the first-stage probabilities at all. Hence one cannot then *define* the actuarially-equivalent simple lottery, by construction, since the informational bridge to that calculation has been burnt. If this CE is generated by RDU, then one can apply RDU to evaluate the first-stage lottery using those CE as final outcomes. The Recursive RDU model assumes one set of RDU preference parameters, just applied recursively in this manner. It can be particularly important when evaluating behavior towards insurance products with non-performance contractual risk, illustrated by Harrison and Ng (2018). This is the compound risk that when a claim is submitted the insurance company sends in the lawyers rather than the claims adjusters, as in litigation over claims for Hurricane Katrina.

The evidence for Cumulative Prospect Theory (CPT) is actually very poor in controlled laboratory experiments with financial incentives.[7] When subjects are provided with a house endowment and losses are framed as coming from that endowment, they behave as if the outcomes are defined over the net gain rather than the gross loss. The same qualitative pattern arises if the endowment stake is earned in some manner. When individuals locally "asset integrate" in this manner, there is no role for the sign-dependence that distinguished CPT. It is also remarkable to see how often evidence for RDU over gains is viewed as support for CPT, when it is patently not: see Harrison and Swarthout (2023) for a detailed review and new experiments.

B. Time Preferences

Similarly, for time preferences different theories agree on the definition of the discount factor, as the scalar exchange rate for an individual between a smaller-sooner (SS) amount of money and a larger-later (LL) amount of money, but then decompose it differently. Exponential discounting models attribute all of the discount factor to a constant variable (utility) cost of time delay, where the variability derives

[7] The same is true, for different reasons, for Dual Theory, the special case of RDU that assumes a linear utility function. Dual theory plays a key role in identification of "limited consideration" in behavioral insurance when EUT is not assumed: see Barseghyan and Molinari (2023), for example.

solely from the time horizon between the LL and SS outcomes. Quasi-hyperbolic discounting models in addition attribute some of the discount factor to a fixed (utility) cost of any time delay.[8]

A key feature of the empirical identification of time preferences is to control for the effect of diminishing marginal utility (DMU), defined as usual by the second derivative of the utility function, on the evaluation of LL compared to SS. Since LL is a larger amount than SS, quite apart from the difference in time dating, one would expect that some of the difference in valuation would be associated with DMU rather than entirely due to the time dating. Hence one can undertake some auxiliary experimental task to infer DMU for an individual, by estimating their utility function, and use that to allow inferences about the part of the discount factor due to time dating.[9] This was the insight of Andersen et al. (2008), who used a risk preference task to infer DMU, and then treated that as a "nuisance parameter" to be conditioned out in order to correctly infer discount rates. Nothing in this approach to joint estimation and identifications rests on assuming EUT, or any specific model of discounting behavior. This methodological insight, adding extra tasks as needed to identify nuirance parameters, and then jointly estimating all parameters of interest, is general, and will be used again as we consider the identification of intertemporal risk preferences and subjective beliefs.

C. Intertemporal Risk Preferences as Multi-attribute Risk Preferences

Intertemporal risk preferences are currently modeled in terms of several sharply contrasting structural theories. One imposes intertemporal risk neutrality by assuming an additively separable intertemporal utility function, and formally assumes away any intertemporal risk aversion. This assumption also ties *atemporal* risk preferences and time preferences at the hip, in the sense that they cannot be independent of each other. The other theories allow for some non-additivity, allowing aversion to

---

[8] These are just two of the more important structural models; Andersen et al. (2014) review a wide range of models.

[9] Andreoni and Sprenger (2012) and Laury et al. (2012) agree that DMU must be accounted for, and propose experimental procedures to bypass the need for an extra task. Both procedures entail serious conceptual and empirical issues, and should not be used.

stochastic variability over time or a preference for temporally correlated variability. The specific alternative that we consider to intertemporal risk neutrality only relaxes the additive separability assumption on the intertemporal utility function.[10] It is worth stressing that the non-additivity in question here does not have the "exotic" status that non-additivity has had in the domain of atemporal risk preference. Additive intertemporal utility was only ever a (very) convenient functional form to ease the math.

The concept of intertemporal risk aversion arises by considering how an individual trades off risks that are time-dated. Define a lottery α as a 50:50 mixture of $\{x_t, Y_{t+\tau}\}$ and $\{X_t, y_{t+\tau}\}$, and another lottery β at the other extreme as a 50:50 mixture of $\{x_t, y_{t+\tau}\}$ and $\{X_t, Y_{t+\tau}\}$, where $X > x$ and $Y > y$, and x, X, y and Y are amounts of money. Lottery α is a 50:50 mixture of both bad and good outcomes in time t and $t+\tau$; and β is a 50:50 mixture of only bad outcomes or only good outcomes in the two time periods. These lotteries α and β are defined over all possible "good" and "bad" outcomes x, X, y and Y that satisfy the constraints that $x < X$ and $y < Y$, as well as all possible common mixtures rather than just 50:50. If the individual is indifferent between α and β, we say that she is neutral with respect to intertemporally correlated payoffs in the two time periods. If the individual prefers α to β we say that she is averse to intertemporally correlated payoffs: it is better to have a given chance of being lucky in one of the two periods than to have the same chance of being very unlucky or very lucky in both periods. And if the individual prefers β to α we say that she is attracted to intertemporally correlated payoffs: it is better to have a chance of being very unlucky or very lucky in both periods than to spread the risk over time periods. The intertemporally risk averse individual prefers to have non-extreme payoffs *across* periods, just as the atemporally risk averse individual prefers to have non-extreme payoffs *within* periods.

Again, the principle of identification with the "bottoms up" approach is to design multiple tasks to estimate a structural model. For intertemporal risk aversion one needs a task to estimate atemporal risk preferences, to infer the atemporal utility functions; one needs a task to estimate time preferences, to infer

[10] Epstein and Zin (1989) preferences *require* a specific, empirically-rejected, non-EUT structure on atemporal risk preferences.

the present discounted utility of time-dated outcomes; and one needs a tasks to estimate the manner in which risks over time are traded off, by offering lotteries of the type α and β for varying x, X, y, Y and to infer the potentially non-additive functional form for intertemporal utility. The final inference is conditonal on inferences in the first two stages, and estimation is joint in the sense of evaluating the "full information likelihood" of all three sets of choices. And yet again, none of the steps involves assuming the simplest models of atemporal risk preferences or time preferences. Experiments with this design, using financially motivated choices and subjects representative of the adult population of Denmark, show clear evidence of intertemporal risk aversion (Andersen et al. (2018)).

An obvious reaction is that "this seems like a lot of work to do" just to explain behavioral towards risk over time. The simple response: if not this work, then what are you assuming away that might matter for your inferences? We tally a list of answers to this question in section 4.

The potential importance of this construct for understanding insurance behavior is immediately evident when we try to make sense of the notion of "inertia" in the behavioral literature in section 4, since that is exactly what intertemporal risk aversion is a preference for. But it takes on even deeper significance when we start to consider insurance lines over outcomes with different attributes, assuming that perfect markets do not exist for these attributes. The same economic logic of trading off risky attributes applies with {x, X} referring to health outcomes and {y, Y} referring to money. Or when {x, X} refers to morbidity outcomes for one health condition and {y, Y} refer to morbidity outcomes for another health condition. One can also have or more than two attributes being considered, where we always allow one attribute {t, T} to denote time of receipt of the attribute. And we can allow nested multi-attribute utility structures. In brief, we have a direct way to characterize multivariate risk preferences that do not collapse to one risk preference over W, the fictitious notion of wealth when perfect markets exist between all attributes of interest.[11]

[11] When this fictitious W exists, a necessary condition for there to be *any* demand for an insurance contract is that the cost of hedging with that contract be less than the cost of removing insurable risk with diversification of assets by means of selling shares on any asset claim: see Benston and Smith (1976).

D. Asset Integration

Once we consider a world in which perfect markets do not exist, the rationale for assuming that risk is defined solely over W collapses.

One immediate corollary is that we must consider insurance decisions as just one of many ways that an individual might manage risk. Mayers and Smith (1983) viewed insurance decisions as one part of a consumer's choice problem over a portfolio consisting of marketable assets and non-marketable assets. The former might be viewed as consisting of equity shares in firms, or assets such as cars, houses or land owned by the individual. The latter might be viewed as human capital, including health. They show that

> Sufficient conditions for insurance decisions to be independent of other portfolio decisions are: (1) there is no moral hazard or adverse selection; and (2) the payoffs to the insurance policy are orthogonal to those of all marketable securities, the consumer's gross human capital, and the payoffs to other insurance policies. Although the first restriction is well known, the second has been unrecognized. Moreover, we argue that this omission is not trivial. There are potentially important covariances in the payoffs with other insurance policies and with human capital which lead to different predictions about insurance demands than obtained under the assumption of separability (p. 310).

Although some of the results are in effect the same as the concept of "self-insurance" developed by Ehrlich and Becker (1972), they formally derive from allowing there to be multiple sources of risk facing the individual and the need to make decisions about the individual's complete portfolio. Thus the demand for self-insurance is not driven by one isolated risk, and must be traded off with all other market and non-market risks (Mayers and Smith (1981)).

A second corollary is that we should not be comfortable with relying on risk preferences defined over changes in this fictitious W. What may be a "reasonable" level of risk aversion over this fictitious W may be very different than the "reasonable" level of risk aversion over the components of W that are relevant for insurance decisions. There have been some modern confusions on this matter, but they are easily sorted out when one tests the theory correctly.[12]

[12] The reference here is to the "calibration critique of EUT" by Hansson (1988) and, much later, Rabin (2000). The claim is that small-stakes risk aversion, as allegedly seen in laboratory experiments, leads to implausible large-stakes risk aversion. One minor aspect of these claims is the reliance on bounded utility functions, such as Constant Absolute Risk Aversion specifications. But there are deeper issues, and Cox and Sadiraj (2006) proposed

E. Subjective Beliefs

A central issue for identification in many of the "tops down" applications of behavioral insurance is the subjective belief that the insured has with respect to expected claims. In general, there are several events that there could be beliefs over. One is whether a loss occurs at all, a binary event. Another is the size of the loss, conditional on it occurring, typically a continuous event.[13] And yet another is the risk of non-performance of the insurance contract, an event that could be binary, discrete or continuous.

Does anything change in the analysis when we explicitly allow for subjective beliefs in the evaluation of an insurance policy? As often with economics, the answer is yes and no. Nothing changes if we assume, following Savage (1972), that decisions are made as if one obeys the ROCL axiom. But things change radically if one does not make that assumption. This seemingly technical issue has great significance for the evaluation of insurance behavior.

For example, consider the subjective beliefs embodied in the forecast cones for the path of Hurricane Ian from 2022 in Florida, and then characterize the forecast cones as just "risks of risks." To simplify, let there be just three possible probabilities that Ian will strike Fort Myers: 0.6, 0.7 and 0.8. Then think of the confidence in each of those probabilities as itself another source of risk, the risk that the forecast is correct.

---

an elegant design to implement a test of this claim, building on the ability to vary "lab wealth" for a given subject, as required from the formal premisses of the claim. Evidence from university undergraduates in the U.S. indicates that the premise is simply false for that population (Harrison et al. (2017)), although evidence from representatives of the adult Danish population shows that the premise is valid for the range of lab wealth considered (Andersen et al. (2018)). In the latter case there are alternative assumptions about the degree of asset integration between field wealth and lottery prizes that allow the reconciliation of small stakes risk aversion with plausible high stakes risk aversion under EUT, and these assumptions appear to apply to the Danish population.

[13] It is common in certain fields, such as health economics, to keep these two events separate by using "hurdle models" (e.g., Coller et al. (2002)). These are econometric specifications that posit some model, such as probit, for the binary event, and some model, such as constrained Ordinary Least Squares, for the conditional loss amount. The claim distribution is then formally modeled as a mixture of these two. One reason for keeping the two events separate is that the data-generating processes driving the two are often very different: the factors leading someone to clinically present to a doctor or hospital for care (e.g., poverty) need not be the factors leading someone to cost more to attend to after they clinically present (e.g., obesity). Hurdle models are called for when the data show a characteristic "spike" at zero, clearly distinct to the eye from the conditional positive values. This is common with insurance claims data, of course. It is then statistically incorrect, but common, to use "tobit" specifications.

What if we have the same level of confidence in each of these three loss probabilities, and completely rule out lower or higher probabilities? Then our subjective beliefs can be summarised in the display in the top left corner of Figure 1. There we see that each of the loss probabilities on the horizontal axis has a confidence, shown on the vertical axis, of 33⅓%. So our confidence that Fort Myers will be hit with probability 0.6, or 0.7, or 0.8 is 100%, but we cannot be more precise than that. It is apparent that the weighted average probability in this case is $0.7 = (⅓ \times 0.6) + (⅓ \times 0.7) + (⅓ \times 0.8)$. This weighted average is shown in the top left panel of Figure 1 by the vertical, dashed line.

Now we revise the forecasting model, much as the National Hurricane Center was doing in real time as Hurricane Ian evolved, and we become slightly more confident that the loss probability is 0.7 rather than 0.6 or 0.8. Then our subjective beliefs might be like those in the top right corner of Figure 1. Here we assume that we now only have 30% confidence that the loss probabilities are 0.6 or 0.8, but we have 40% confidence that the loss probability is 0.7, so the confidence-weighted average is again 0.7. As the hurricane progresses and tightens in on Fort Myers, and the forecasting model is updated, the confidence we attach to the loss probability being 0.7 gets larger and larger in the bottom two panels of Figure 1. This example has been constructed so that the confidence-weighted average in all four cases remains the same, 0.7.

Where do the belief distributions in Figure 1 come from, when we turn to empirical identification? Again, we need an additional task to *observe reports* about the events, from which we can *infer beliefs* abut the event. These tasks take the form of scoring rules, the most popular of which is the Quadratic Scoring Rule (QSR): see Harrison et al. (2017).[14] And yet again, these procedures for eliciting subjective beliefs do not assume EUT or SEU. We do need to distinguish *conceptually* between someone having subjective probabilities from whether they act "optimistically or pessimistically" towards those (subjective or

[14] If the subject is risk-neutral, these QSR reports can be taken directly as beliefs. If not, one can infer them making some conditional assumptions about risk preferences, as demonstrated by Andersen et al. (2014) and Harrison et al. (2022), or by extending the QSR to be defined over binary rewards that "risk neutralize" the agent, as in Harrison et al. (2014) and Harrison et al. (2015b).

objective) probabilities, and the empirical tools allow us to do that. Subjective beliefs about loss probabilities and claim amounts are a challenging confound to many field inferences about insurance, whether or not an experiment was conducted.

Figure 2 shows an asymmetric subjective belief, with 10% confidence in the probability π being 0.6, 60% confidence in the probability π being 0.7, and 30% confidence in the probability π being 0.8. Now consider a lottery in which one gets \$X if the event occurs, and \$x otherwise. Then the Subjective Expected Utility (SEU) of this lottery is

$$0.1\times0.6\times U(X)+0.1\times0.4\times U(x)+0.6\times0.7\times U(X)+0.6\times0.3\times U(x)+0.3\times0.8\times U(X)+0.3\times0.2\times U(x),$$

which collapses to

$$(0.1\times0.6 + 0.6\times0.7 + 0.3\times0.8) \times U(X) + (0.1\times0.4 + 0.6\times0.3 + 0.3\times0.2) \times U(x)$$

and hence to $0.72 \times U(X) + 0.28 \times U(x)$ under ROCL. So the non-degenerate distribution in Figure 2 can be boiled down to a degenerate subjective probability π of 0.72 under ROCL: an impressive identifying restriction. In summary, we can define risk as existing whenever there is objective risk *or* when there are subjective belief distributions to which the individual applies ROCL.

We can now see the formal difference between risk and uncertainty. Looking at the four levels of confidence in Figure 1, would you treat these the same if you lived in Fort Myers and were deciding how to manage the risk of the hurricane? Surely the same logic that suggests that you might be averse to simple risks suggests that one might also be averse to your own lack of confidence in the top left panel? This is sharpest when you consider the two extremes of confidence, where the 0.7 loss probability goes from having 33⅓% confidence to having 90% confidence. If this response is the case, and the *shape* of the subjective belief distribution, *apart from the average*, matters for your risk management choices, we say that you are facing uncertainty and that you might be averse to that uncertainty. This is just like risk aversion over simple, direct risks, but arises because you do not treat compound risks, here the confidence of your beliefs, the same way as those simple risks, as required by ROCL. Uncertainty aversion arises when you are averse to the imprecision of your own beliefs, as in the top left corner, and would be willing to pay to have more confidence. In other words, uncertainty arises when you have a subjective probability distribution of beliefs but do not apply ROCL to it.

How we relax ROCL is a matter for important, foundational research. One popular approach is the "smooth model" of Klibanoff et al. (2005), with important parallels in Nau (2006) and Neilsen (2010). We can illustrate the smooth model with a simple example. Let CE(π=0.6) be the Certainty Equivalent (CE) of the lottery $0.6\times U(X)+0.4\times U(x)$, CE($\pi$=0.7) be the CE of the lottery $0.7\times U(X)+0.3\times U(x)$, and CE($\pi$=0.8) be the CE of the lottery $0.8\times U(X)+0.2\times U(x)$. Then the evaluation of the lottery can be written

$$0.1\times\varphi(CE(\pi=0.6)) + 0.6\times\varphi(CE(\pi=0.7))+ 0.3\times\varphi(CE(\pi=0.8)),$$

where φ is a function defined over the CE of the lottery that is conditional on a particular subjective probability value. Akin to the properties of $U(\cdot)$ defining risk attitudes under EUT or SEU, the properties of $\varphi(\cdot)$ define attitudes towards the uncertainty over the particular subjective probability value.[15] If φ is concave, then the decision-maker is uncertainty averse; if φ is convex, then the decision-maker is uncertainty loving; and if φ is linear, then the decision-maker is uncertainty neutral. The familiar SEU specification emerges if φ is linear, since ROCL then applies after some irrelevant normalization. The overall evaluation of the lottery depends on risk attitudes *and* uncertainty attitudes, and there is no reason for the decision-maker to be averse to both at the same time. An important econometric corollary is that one cannot infer attitudes toward uncertainty from observed choice until attitudes toward risk are characterized.

We can go one step further, and consider the situation in which you know *something* about the risks, but not *everything*. Still assume you are confident that the possible loss probabilities are only 0.6, 0.7 or 0.8 as before. But now consider the implications of you only knowing that you have 33⅓% confidence

[15] In the original specifications φ is said to characterize attitudes towards ambiguity, to be defined momentarily. Schmeidler (1989; p.582) explicitly characterized such specifications as uncertainty aversion, noting that "Intuitively, uncertainty aversion means that 'smoothing' or averaging utility distributions makes the decision maker better off."

in the 0.6 loss probability. Figure 3 displays four possible cases, out of a huge number of possibilities. Every one of these possible cases assigns 33⅓% confidence to the 0.6 loss probability, but assigns the remaining 66⅔% confidence differently. We display the known confidence in a dark shade in Figure 3, and the unknown, possible confidences for loss probabilities 0.7 and 0.8 in light shade. In this situation we say that you face *ambiguity* about risks, following Ellsberg (1961), since there are many ways in which the things you do not know might get resolved – we just show four of them here.

There is considerable debate among economists as to how people respond to ambiguity, as well as debate over recommendations about how they *should* respond to it. Some might suggest that you just fill in the blanks with equal weight, on the basis of the alleged "principle of insufficient reason," leading to the well-defined subjective belief distribution in the top left corner of Figure 3. You know these two outcomes, loss probabilities of 0.7 and 0.8, jointly have a weight of 66⅔% in your assessment, but you have no informed basis to allocate that overall weight between the two. Others suggest that you might exhibit *ambiguity aversion* by assuming the worst possible outcome, following Gilboa and Schmeidler (1989).[16] If you live in Fort Myers, this means you would assign the full 66⅔% to the loss probability of 0.8, as shown in the bottom left corner of Figure 3. Of course, once you have assigned the other weights in one of these or other ways, you may or may not exhibit uncertainty aversion as well.[17]

4. Selected Topics in Behavioral Insurance

The goal of this selective review is to examine important and influential contributions to behavioral insurance, with an eye to seeing why they are valuable contributions, and also what remains to

[16] The decision rule to assume the worst possible outcome as a response to ambiguity is not, as often interpreted, just assuming deep pessimism in the face of ambiguity. Instead, it is just showing formally how one could still have *some* decision rule in such a low-oxygen informational setting.

[17] What if you had the extreme case of knowing that the probabilities are 0.6, 0.7 or 0.8, but not even knowing the confidence for the 0.6 loss probability? Then again, the first decision rules just outlined would tell you to assign 33⅓% to each loss probability, as in the top left panels of Figures 1 and 3. And the second decision rule would tell you to assign 100% to the worst outcome, and behave as if you are completely confident that the loss probability is 0.8.

be done.

A. Perfectly Informed Risk Types?

A common identifying assumption in many behavioral studies of insurance and annuity choice is that individuals know their own risk type. Moreover, it is further assumed that it happens to be *exactly* the risk type that the actuaries at an insurance firm might infer.

Cohen and Einav (2007) examine a rich data-set of choices over menus of deductibles and premium payments for auto insurance that varied across individuals. These menu options constitute necessary controls to view these data as a natural field experiment. The researchers know the premium offered, but do not know the subjective perception of the risk of a claim, or the risk that the claim will be paid in full. To proxy the latter they assume that individuals have accurate *point* estimates of the true distribution, a tenuous assumption even for experienced drivers. Moreover, they must assume EUT, since they have no way to identify non-EUT models of risk preferences, and hence the calibration implications of such preferences. Certain non-EUT models of risk preferences, such as RDU, have been shown to dramatically affect the valuation of insurance when calibrated to estimates from real choices in the field: see Hansen et al. (2016).

The key identifying assumption, that individuals know the actuarial loss rates and claim values, turns out to play a critical role in most of the observational literature as well. In a survey Ericson and Sydnor (2017; p.54) correctly note that, "When economists analyze health insurance markets, they typically assume that people are aware of the distribution of their possible medical bills for the year and choose their health plan with that information in mind." In fact, most studies go well beyond assuming awareness of the *distribution*, and are assumed to have statistically degenerate beliefs on some *scalar statistic* derived from that distribution.

We noted earlier that assuming that an individual makes decisions over risky outcomes by reacting optimistically or pessimistically to objective risks is *not* the same as assuming that individuals might have subjective perceptions of risk that deviate from objective risks. Of course, the two might be impossible to tease apart in field settings, but it is easy to do in theory and controlled laboratory experiments that operationalize that theory. The implications of teasing these apart are apparent when one starts to engage in normative tinkering: one might plausibly adopt a different normative stance towards subjective beliefs being different from the beliefs of some actuaries than the stance one takes towards optimism or pessimism with respect to those subjective beliefs.

We can see the difficulty that RDU poses for inference about insurance choice when one allows for subjective probabilities in Barseghyan et al. (2013). They exploit the fact that the decision-makers in their sample had a choice from multiple deductibles, and recognize that this allows them to identify the role of diminishing marginal utility *and* "probability weighting" in the sense of RDU, since these two channels for a risk premium have different implications at different deductible levels. They also explicitly acknowledge that what they call probability weighting might also be simply subjective risk perceptions that differ from the true claims rate, noting that their analysis "does not enable us to say whether households are engaging in probability weighting *per se* or whether their subjective beliefs about risk simply do not correspond to the objective probabilities" (p. 2527). Their striking result is that probability overweighting (or, we add, subjective risk bias) with respect to claims is, along with diminishing marginal utility, a central determinant of the risk preferences of these deductible choices.

A critical assumption that they make, common to most of the studies of observational data, is to estimate a *scalar* loss probability for each individual or household in their data. To be sure, these estimates invariably use a rich dataset of demographic characteristics from the data, and presumably accessible to the actuaries and underwriters of the insuirance contract. So they have that level of credibility. But in all cases a point estimate is assumed as if known by the decision-maker, not some subjective probability distribution around that point estimate. To be specific, this assumption is used in Barseghyan et al. (2013; p. 2505)(1997; p. 1997)(2021; p. 2028) and in Barseghyan and Molinari (2023; p. 1021). It also plays a key role in the evaluation of health insurance in the Netherlands by Handel et al. (2020; p.11ff.). The sole exception appears to be Handel et al. (2019).

B. "Inertia"?

Most insurance contracts have limited contract horizons, usually one year, and are then renewed with potentially different coverage and premia offerings. The behavioral literature often just states that those purchasing insurance exhibit "inertia," implying that they mindlessly renew contracts even when there appear to be better alternatives available. Of course, this behavior sounds exactly like intertemporal risk aversion, introduced in section 3.

Handel (2013) exploits a natural field experiment in which a large firm changed health insurance options from an active choice mode to a passive mode in which the previously selected choice was the default choice in later years unless action was taken. This change allowed inferences about the role of "inertia" in insurance plan choice. The behavior of new employees, who needed to make an active choice when previous employees were faced with passive choices, provides intuition for the significance of inertia, assuming comparability of other characteristics between the two employee groups. Some passive employees faced "dominated" choices over time as insurance parameters changed, and their sluggishness in the face of these incentives provides indicators of inertia; the use of scare quotes around the term word "dominated" will be explained momentarily. Risk preferences are assumed to be distributed randomly over the population sampled, and be consistent with EUT. Individuals know their own risk preferences, but this is unobserved by the analyst. This could cause identification problems if the "nonfinancial attributes," to use the expression of Handel and Kolstad (2013), also varied across all plan choices, but three Preferred Provider Option (PPO) plans had no differences in these attributes: hence their variations in "financial attributes," such as deductible, coinsurance, and out-of-pocket maxima, could be used to identify (atemporal) risk preferences. In keeping with other observational studies, the distribution of claims was simulated using sophisticated models akin to how an actuary would undertake the task, and individuals were assumed to know the risks they faced exactly.

Since the focus is on "inertia" over time, a critical and implicit behavioral assumption is that individuals are *intertemporally risk neutral* with respect to the attributes of the health plan over time.[18] An individual that is *intertemporally* risk averse cares, as a matter of preference, that attributes not vary *over time*. If individuals are assumed to be intertemporally risk neutral then they do not care about variations in attributes over time, as one moves from plan to plan over time, as long as the average attribute remains the same.[19] So giving up their favorite family doctor for a new family doctor does not matter at all, *ceteris paribus* the average attributes of the doctor, and will be accepted willingly for any tiny improvement in premia. For now, assume that the sole attribute considered is the time spent with the doctor, not the identity of the doctor or whether one has a history with the doctor. Then it is being assumed that this non-financial attribute is the same on average, and the plans can be viewed as dominated on the basis of the financial attributes. The focus here is on an oft-mentioned attribute that, as a matter of fact, was the same across the PPO plans that the individuals being studied could choose from.

But it is clear, as emphasized by Handel and Kolstad (2013; p. 2451) that there is evidence that 50% of subjects did not think that the attributes were the same across the PPO plans or were not sure of it. All that is needed is that individuals do not *subjectively* believe that these attributes are the same across these PPO plans. This is a false subjective belief: it is not a friction. Given this false belief, the preference for not changing plans could be due to a preference for stability of attributes over time, which is what intertemporal risk aversion is all about. Given this false belief, the plans are not subjectively dominated in terms of the financial attributes. Given this false belief, what is attributed to "inertia" is exactly what a preference for temporal stability implies when one allows for it. And the methodological point is more

[18] In general the reference to attributes should include what are referred to as "financial attributes" as well as "nonfinancial attributes," but in the context of Handel and Kolstad (2013) the term just refers to the latter. And -for present purposes the formal theories of multi-attribute risk aversion can be viewed as including intertemporal risk aversion as a special case, where one of the attributes is whether the attribute is consumed sooner or later. Similarly, one can define multi-attribute risk aversion even if there is no time-dating of outcomes.

[19] This is separate from the assumption that "consumers are myopic and do not make dynamic decisions whereby current choices would take into account inertia in future periods" (p.2662). That assumption has to do with sophistication with respect to the effect of current consumption on future consumption, akin to "rational addiction" models.

general, of course, when we consider plan choice over options with objective differences in attributes.

In the context of the data evaluated by Handel (2013), intertemporal risk aversion is just a taste for *not* having variability in claims risks over time, where risks refer to all subjective financial and non-financial attributes of the plan, and that is met simply by choosing the same plan year over year. Just as one is willing to pay a risk premium in terms of expected value to reduce atemporal risk aversion, the willingness to put up with lower expected value plans can be seen as a risk premium to reduce intertemporal risk aversion with respect to attributes.[20]

C. Risk Preferences Versus Information Frictions?

Handel and Kolstad (2013) seek to tell a story about the role played by "risk preferences" and the role played by "information frictions" in determining the demand for health insurance products. They also seek to tell a story about the welfare implications of the inclusion of "information frictions." I use the expression "seek to tell a story" to be clear that this is academic rhetoric, for the purpose of shifting discussion away from just assuming that "risk preferences" alone explain insurance behavior.[21] Others might not see this type of rhetoric as the right way to model behavior, but that position neglects any appreciation of the paucity of data with which to draw inferences in the field.

Handel and Kolstad (2013) start with a rich administrative data set in which individuals with certain demographic characteristics had to choose between two health insurance plans. One PPO plan

---

[20] This point has fundamental implications for the resulting welfare analysis (p.2669-2679). The story here is that "consumers enroll in sub-optimal health plans over time, from their perspective, because of inertia. After initially making informed decisions, consumers don't perfectly adjust their choices over time in response to changes to the market environment (e.g., prices) and their own health statuses" (p.2669). Another story, equally consistent with the observed choices and EUT, is that consumers have a preference for avoiding subjective intertemporal risk in the health plan lotteries they choose. And yet another story has to do with where the false beliefs came from, in this specific context.

[21] They reference (p. 2450) Cohen and Einav (2007) and Bundorf et al. (2012) as conducting welfare analysis of health insurance plans in which they use "observed choices to identify risk preferences." In fact, risk preferences are not identified by Bundorf et al. (2012). And Cohen and Einav (2007) undertake no welfare analysis. Similarly, Einav et al. (2010a; p. 878) claim that Einav et al. (2010b) and Bundorf et al. (2012) "recover the underlying (privately known) information about risk and preferences." Neither of these are true.

provides “comprehensive risk protection” (p. 2451); the other plan, a High Deductible Health Plan (HDHP), provided access to “the same medical providers and treatments as the PPO, lower relative upfront premiums, and larger relative risk exposure.” (p. 2451). In addition to the administrative data, for a significant sub-sample of the population they also had a linked survey of beliefs about these plans. The intuition of their results can be seen by one example (p. 2451): if 50% of individuals incorrectly believed that the PPO provided greater medical access to providers and treatments (20%), or were not sure about that (30%), they were more likely to choose the PPO than individuals that knew that the plans provided the same access. Call these subjective beliefs about some core attributes of the products. Given these subjective beliefs, apply SEU to these choices, and what we see is just a better apple or a less risky apple being selected over a poor apple. The first 20% subjectively perceive a more useful product, and the second 30% subjectively perceive a less risky product.

The first formal step in the analysis is just to recover risk preferences from observed choices between the PPO and HDHP. In this case the model assumes EUT, and critically *assumes that individuals know the actuarial probabilities* of receiving benefits from each insurance plan. Intuitively, think of the PPO as the safe lottery and the HDHP as the risky lottery.[22] To borrow an expression, the resulting estimates of risk aversion are “just wild and crazy guys,” to be laughed at because they are so high (p. 2452). Of course, we know from RDU models of risk preferences that this *might* actually be a combination of (very) pessimistic beliefs about receiving the benefits of the HDHP and a (modestly) concave utility function. The point is that the available data is unable to differentiate them, hence we cannot claim to have identified risk preferences without accepting the maintained assumption of EUT for all individuals, and where EUT assumes remarkably prescient knowledge of the actuarial risks of what are clearly compound subjective lotteries.

[22] The effort to construct these actuarial probabilities (p. 2480) is impressive. It uses *ex post* information to predict the utilization of four types of health expenditure in the coming year, and then *ex post* data on the costs of each of these expenditure types to predict spending distributions. One could use these objective calculations as the basis for eliciting subjective probability distributions with incentive-compatible experiments, which is what we need to estimate an SEU model of insurance choice.

The second formal step in the analysis is to correctly recognize (p. 2455ff.) that modern health insurance plans have many attributes that differentiate them. We are not in a world, at least for these product lines, of just trading off lower deductibles for higher premia. In the absence of these "nonfinancial attributes" the utility function has, as an argument, W_k - P_{kj} - s_i where W_k is wealth for household k, P_{kj} is the premium that household k faces for insurance plan j, and s_i is the out-of-pocket payments for some sad event i. Then there is some actuarial probability mass function, let us assume, defined over the s_i, and that depends on the household k and plan j in question. Now consider the effect of "nonfinancial attributes," such as "the network of physicians and hospitals available, the time and hassle costs associated with dealing with claims, and the tax benefits of linked financial accounts." (p. 2455). For short, call this $BLOB_j$ for plan j, recognizing that BLOB has potentially many arguments reflecting a vector of perceived attributes.[23] The argument of the utility function then becomes W_k - P_{kj} - s_i + $BLOB_j$. This specification is at the heart of the analysis.

A theoretical problem with this way of handling "nonattribute frictions" is that they are included in an additive manner. This implies that they are known quantities if one knows the household k and plan j, so they are not themselves risky.[24] This further implies that even if they were assumed to be risky, they *cannot* trade off with other "financial risks." The general point is that we are talking about "risk preferences" here, albeit in the form of an exciting cocktail of multi-attribute risk preferences, but just risk preferences nonetheless.[25]

The modeling upshot is that I am suggesting a different "story" here, and there is no possible way for these data, as rich as they are in comparison to most observational data sets, to tell them apart. But this

---

[23] Indeed, BLOB could be viewed as a nested utility function defined over these attributes, as proposed in footnote 12 (p. 2456) and in the empirical model. In the empirical model (p. 2475) these attributes are all treated as binary, and included additively.

[24] The only stochastic aspects of these attributes (p. 2456) is that they are *observed* with error by the researcher, reflecting unobserved but deterministic heterogeneity.

[25] Handel and Kolstad (2015; p.2452) include "inertia" in their structural model, and comment that "incorporating inertia into the model matters a lot for risk preference estimates." They refer here to *atemporal* risk preferences. The deeper implications for risk preferences, having to do with *intertemporal* risk preferences, is discussed earlier with reference to Handel (2015), where "inertia" is the main story.

story has very different implications for how one does descriptive and normative evaluations of observed insurance choices.

D. Subjective Risk Perceptions, Perhaps?

Einav et al. (2010b) develop a structural empirical model of the demand for annuities in the United Kingdom between 1988 and 1994 for which the annuitant was still alive at the start of 1998. Data on gender, age at annuitization, and age at death if prior to 2006, is observed, as well as the level of annuitization and the choice of a 0, 5 or 10-year guarantee. Annuitization itself is compulsory for most of the accumulated balances from tax-preferred, defined-contribution pension payments. Annuity payment rates decline with longer guarantee periods (Table II, p.1039) and this pattern was held constant over the period of annuitizations.

A key issue for the effects of adverse selection on welfare evaluation is whether there is any private idiosyncratic information that individuals have when they decide on the length of guarantee. Subjective beliefs about longevity, conditional on reaching the age at which this decision is made, are what is relevant for *ex ante* welfare evaluation. However, *ex post* mortality rates can provide some partial indicator of the potential extent of the problem. Over all 9,364 annuitants, 10%, 87% and 3% chose the 0-year, 5-year and 10-year guarantee, respectively (Table I, p. 1037). Conditional on choosing the 0, 5 or 10 year guarantee, mortality rates were 16%, 21% or 19%, respectively. Across the three contracts to choose from, the mortality rates were 20%, so 1-in-5 received the *ex post* benefit of the guarantee. Of course, this can only be one piece of the puzzle: subjective beliefs about these mortality rates, even if they are assumed to match the realized rates, do not tell use subjective beliefs about longevity *beyond* the guarantee period.

Heroic assumptions are needed to generate welfare estimates of alternative policies for individuals. This is not a criticism, just a recognition that if one is to go beyond the qualitative identification of the existence of adverse selection or moral hazard in insurance purchases and quantify their importance, one must let theory, parametric structure and assumptions play a central role. In this case the empirical model assumed EUT decision makers, Exponential discount rates, and additive intertemporal utility functions (p.1041): these assumptions rule out alternative, popular behavioral assumptions which could be important for many individuals.

Furthermore, the same CRRA utility function over consumption applies to all individuals, and with one caveat the same CRRA utility function used for consumption applies to all individuals with respect to the utility of bequests at death (p.1043).[26] Values for RRA and the discount rate are assumed, *not estimated*. In fact, RRA is set to 3, and the discount rate is set to 4.3% p.a. on the basis of a real interest rate at the beginning of 1992. In addition, since annuitization rates are in nominal currency units, they have to assume an expected annual inflation rate of 5% to infer the real annuity payout stream that individuals are choosing over.[27] Although there are several references to "estimates of the joint distribution of risk and preferences" (e.g., p.1082), there is no sense at all in which risk preferences towards consumption variability or bequest risk are estimated, let alone time preferences, let alone any interaction between risk and time preferences.[28]

Some of these heroic assumptions are treated as being relatively unimportant. Perhaps the most important for this insurance contract is longevity risk, as noted earlier. Einav et al. (2010b; p. 1079) claim that, "Throughout we made a strong assumption that individuals have perfect information about their actual mortality rate [...]. This is consistent with empirical evidence that individuals' perceptions about their mortality probabilities covary in sensible ways with known risk factors, such as age, gender, smoking,

---

[26] The caveat is that individuals have a multiplicative weight that they put on the latter argument of utility, interpreted as "the relative weight that individual *i* puts on wealth when dead relative to consumption when alive" (p.1043).

[27] Hansen et al. (2016) illustrate how one can combine estimates of key behavioral parameters such as these, from field experiments with representative populations, and insurance data from the same population.

[28] The grand utility function over consumption flows while alive and bequest motives when dead is also assumed to be additive in these two components, quite apart from the additivity of the former over remaining time periods of life. Inflation risk is also assumed away (footnote 11, p.1050). The point of mentioning all of these factors is that tradeoffs between these many risks are ruled out by additive structures. The implications for behavior towards retirement planning, and related life-cycle decisions, of allowing tradeoffs between risk preferences and longevity risk in intertemporal settings has been extensively explored by Bommier (2006)(2010)(2013) and Bommier and Rochet (2006).

and health status [...] Of course, such work does not preclude the possibility that individuals also make some form of error in forecasting their mortality." In fact, there is evidence that individuals, as well as epidemiologists and actuaries, have significant struggles with forecasts of longevity risk: see Elder (2013) and Di Girolamo et al. (2015).

In addition, and again a common assumption when working with observational data, individuals are assumed to know *exactly* the relevant risks that are being insured, in this case their own longevity risk (p.1042). As explained in section 3, we can view this as a "degenerate" subjective risk if we assume SEU, which is to say that we assume that individuals have subjective belief distributions about their longevity risk *and apply ROCL* to reduce them to their weighted average when making annuity decisions. Important extensions would consider the effects of uncertainty aversion or ambiguity aversion.

E. Insurance Literacy?

Loewenstein et al. (2013) report the results of hypothetical surveys to evaluate if individuals understand the health insurance products they are being asked to purchase.[29] One survey asked about some basic insurance concepts (deductible, copay, coinsurance, and out-of-pocket maximum), and then presented a standard, commercial health insurance contract with all of these concepts in play and asked some questions about what the contract entailed. Accepting the methods to measure insurance literacy for the moment,[30] the conclusion is that there is "strong evidence that consumers do not understand

[29] Ericson and Sydnor (2017; p.58ff.) review the broader literature on "confusion" in health insurance choice.

[30] Some of the questions are not ideal measures of literacy in this domain, reflecting poor survey design for the inferences intended. For example, the questions about the concept of deductibles has multiple-choice answers (p. 853) to the question "Which of the following best describes a Deductible?", and responses are coded as "true or false." Two of the responses are clearly false, one is "I'm not sure," and the two others are "The amount you pay before your insurance company pays benefits" and "The amount you pay before your health expenses are covered in full." The last one is false in the sense that it ignores possible co-pays that might apply over the period covered, and ignores possible co-insurance payments. But it certainly covers the essential idea reasonably well. Another question poses a specific scenario about the commercial plan, Plan T on page 861, for which the correct answer requires the arithmetic evaluation of $\$1,500 + 0.8 \times (\$100,000 - \$1,500) = \$1,500 + 0.8 \times \$98,500$. The respondent is asked for an open-ended response: do you know the exact answer without calculating it?

traditional plans" (p. 850). So how do we evaluate this? We are told that "limited understanding is likely to lead to suboptimal decisions," (p. 852), but how do we know? If someone responds to a survey question "I'm not sure," that is a *plausible* signal for someone that is likely to seek a cognitive scaffold prior to making an actual decision (e.g., check the internet, check with an expert, or just check with a friend). Access to a scaffold does not ensure an optimal decision, but the response is at least flagging *some* lack of confidence in the answer, and that surely has some implications for behavior beyond just assuming *a priori* that someone will pick at random or in systematic error.

Loewenstein et al. (2013) do flag two further ways in which understanding, or literacy as it should be termed, might affect individual welfare. One is whether individuals choose health insurance policies that minimize their expected costs. This is a problematic metric, hinted at with the comment that "while cost minimization is not necessarily equivalent to utility maximization, it is a useful benchmark." (p. 852). A more accurate statement would be that "cost minimization is not equivalent to expected utility maximization, or even maximization of some other interesting utility function, and is not a useful benchmark." We simply have to minimally attend to risk preferences, time preferences and subjective beliefs before we start making claims about individual welfare. The second way in which literacy failings might impact individual welfare is to see if a "lack of understanding was correlated with their insurance choices," as in Handel and Kolstad (2013). In the absence of more nuanced evaluations of these choices, in terms of preferences and beliefs, such correlations mean little.

Literacy is a deeper issue, and one that has been neglected in behavioral analyses of insurance demand. Fang et al. (2008) stressed the importance of cognitive ability as a determinant of demand for insurance by older consumers, and the specifics of their survey data point to their measure detecting correlates of literacy. Fang et al. (2006; p.17) explain:

> Our measure of a person's cognition combines his/her performance on four different tests/questions: word recall, a Telephone Interview for Cognitive Status (TICS) score, subtraction, and numeracy. These scores may proxy for an individual's degree of economic "rationality," i.e. his/her ability to think through the costs and benefits of Medigap insurance. There is a large body of literature showing that many of the elderly have

> difficulty understanding the basic Medicare entitlement, and/or the features of supplemental insurance (see, e.g., Harris and Keane (1999) for empirical evidence [...]).

Without endorsing any simple measures of economic rationality, one can easily see that these components of cognition are likely intermediate inputs into the production of literacy with respect to complex insurance contracts.

More formal tests of insurance literacy, tied to specific features of index insurance contracts that allow two-sided non-performance risks, were implemented by Harrison et al. (2022). Subjects in laboratory experiments were given instructions on the index insurance contract they were to make incentivized purchase decisions about, as actuarial parameters were varied exogenously. Prior to making any choices, they were asked to report their beliefs about how different sets of parameters and decisions would affect their possible payouts, showing that they understood that feature of these contracts. These beliefs were elicited with incentivized belief elicitation procedures, providing a rigorous measure of a specific component of insurance literacy that was used to evaluate subsequent insurance purchase choices.

5. Conclusion

The field of behavioral insurance has achieved its initial objective of making sure that rigorous empirical research on insurance behavior and welfare evaluation *consider* alternatives to standard modeling assumptions. It is time now to stop thinking of behavioral insurance as a parallel scholarly universe of slogans, intuitive stories, and "free parameters," and to get down to the harder job of marshaling those rich alternatives in a conceptually correct manner, appropriately for the inferences at hand, and when needed.

Inference from field data, even when those data have been set up naturally to provide some valuable controls, is tenuous because there are so many behavioral moving parts at work when studying insurance behavior. There has been a tendency to find some "magic bullet" from the behavioral repertoire that can account for core patterns of behavior. The rest of the model is stripped down to show the contribution of that novel contribution, which is exactly what good modeling is about when exploring such speculations for the first time. But then the novelty becomes baked in to the modeling as a fact that *must* be accounted for, rather than something we *could* account for, and the severe identifying restrictions just become the price of working with field data. In the words of Homer Simpson (season 5, episode 22), "every time I learn something new, it pushes some old stuff out of my brain."

Normative inference about insurance behavior is also complicated when we relax the axiom of (direct) revealed preference, in order to be able to infer welfare losses rather than just (weak) welfare gains. This challenge arises even if one assumes standard behavioral models of risk preferences and beliefs, and just becomes harder when one also wants to relax those.

We can do better. Doing so means knowing the theory behind the behavioral insights into risk preferences and subjective beliefs, rather than tossing those insights in as an *ad hoc* afterthought in one of the many perfunctory claims of robustness. It is not true that every alternative to EUT does better at explaining behavior. If you explicitly or implicitly assume away any preference for reducing risk over time, do not be surprised if your modeling errors exhibit autocorrelation and look like "inertia." It is not true that one must jump from the extreme of SEU to ambiguity aversion. Nor should we confuse subjective probabilities with optimism or pessimism in relation to (objective or subjective) probabilities, and so on. It also means understanding when one must collect extra data, how to collect that data, and how to use it to condition inferences. And it means thinking of conditioning parameters explicitly as priors, particularly when turning to normative inferences. In turn, the use of priors on conditioning parameters will be greatly assisted by using formal Bayesian methods. There is much remaining to do in behavioral insurance.

Figure 1: Risky or Uncertain Beliefs

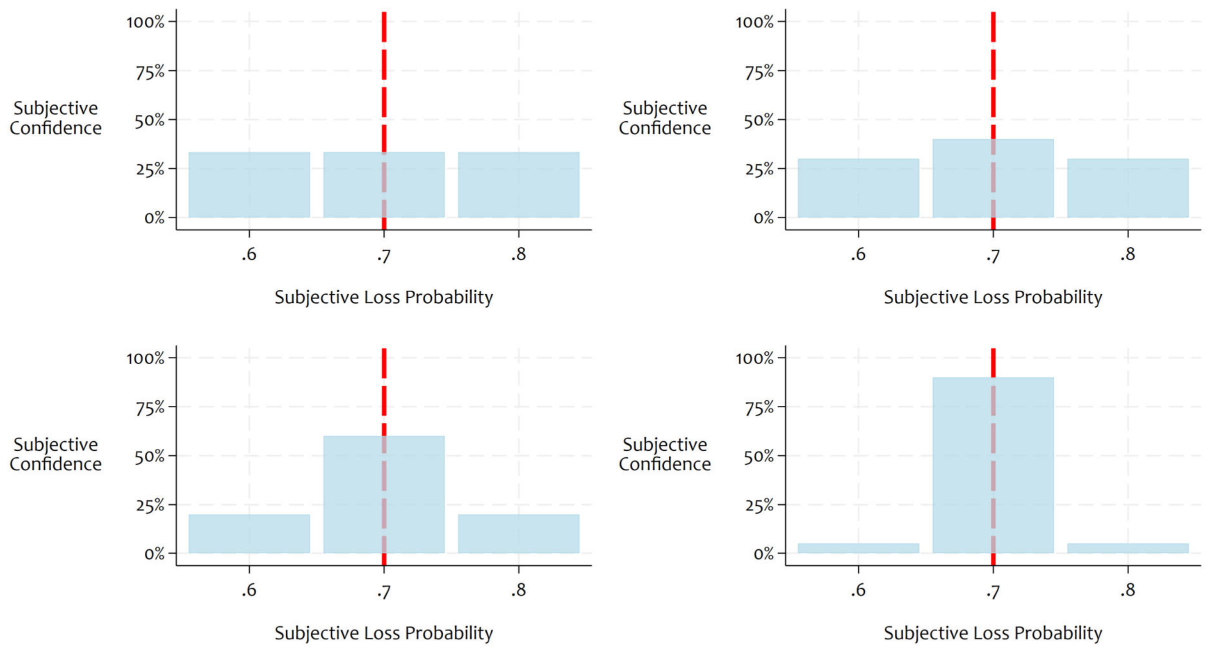

Figure 2: Asymmetric Beliefs

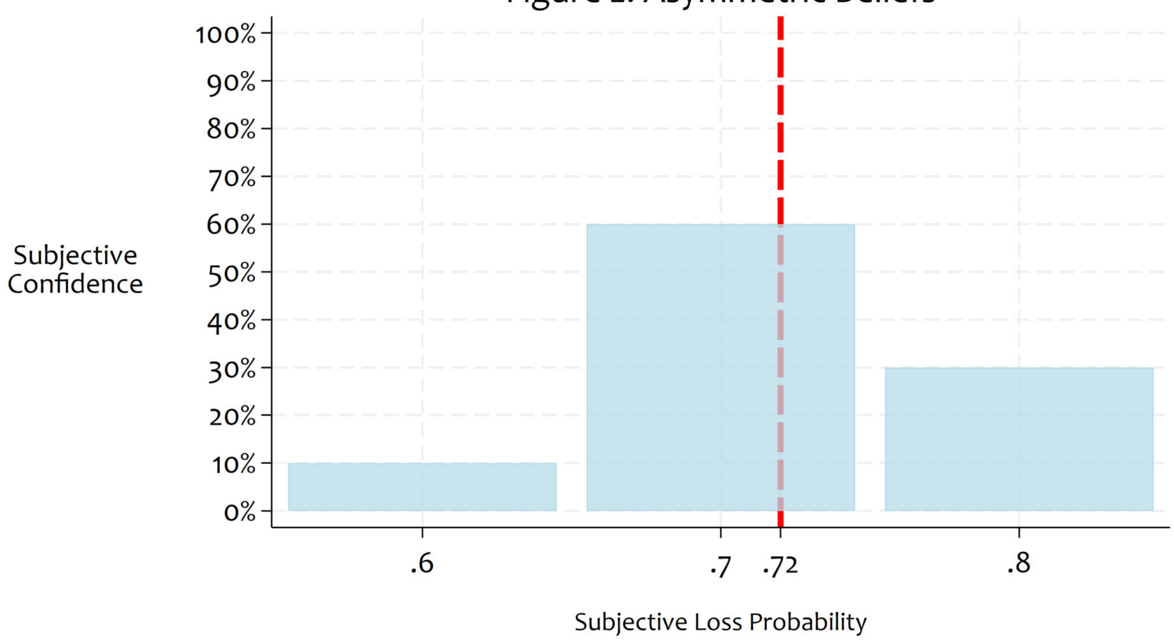

Figure 2: Ambiguous Beliefs

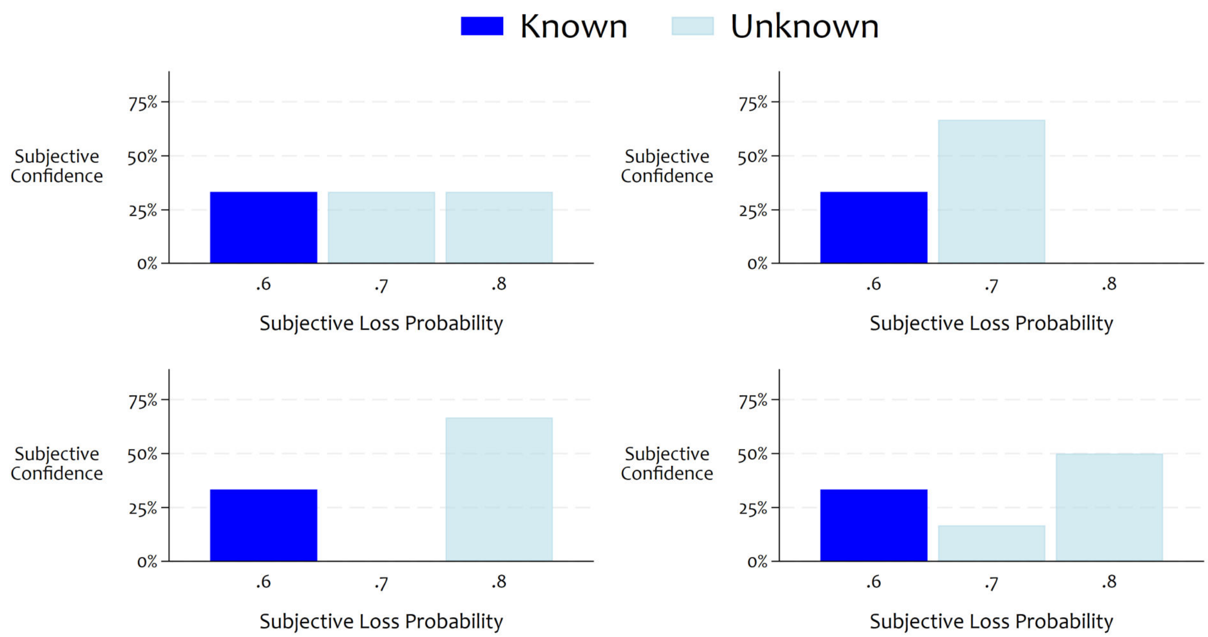

References

Andersen, Steffen; Cox, James C.; Harrison, Glenn W.; Lau, Morten I.; Rutström, E. Elisabet, and Sadiraj, Vjollca (2018) Asset Integration and Attitudes to Risk: Theory and Evidence. Review of Economics & Statistics 100(5): 816-830.

Andersen, Steffen; Fountain, John; Harrison, Glenn W., and Rutström, E. Elisabet (2014) Estimating Subjective Probabilities. Journal of Risk & Uncertainty 48:207-229.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet (2008) Eliciting Risk and Time Preferences. Econometrica 76(3):583–618.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet (2014) Discounting Behavior: A Reconsideration. European Economic Review 71:15–33.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet (2018) Multiattribute Utility Theory, Intertemporal Utility, and Correlation Aversion. International Economic Review 59(2):537–555.

Andreoni, James, and Sprenger, Charles (2012) Estimating Time Preferences from Convex Budgets. American Economic Review 102(7):3333-3356.

Barseghyan, Levon; Coughlin, Maura; Molinari, Francesca, and Teitelbaum, Joshua C. (2021) Heterogeneous Choice Sets and Preferences. Econometrica 89:2015–2048.

Barseghyan, Levon and Molinari, Francesca (2023) Risk Preference Types, Limited Consideration, and Welfare. Journal of Business & Economic Statistics 41(4):1011-1029.

Barseghyan, Levon; Molinari, Francesca; O'Donoghue, Ted, and Teitelbaum, Joshua C. (2013) The Nature of Risk Preferences: Evidence from Insurance Choices. American Economic Review 103(6):2499-2529.

Barseghyan, Levon; Molinari; Francesca, and Thirkettle, Matthew (2021) Discrete Choice under Risk with Limited Consideration. American Economic Review 111(6):1972–2006.

Benston, George J., and Smith, Clifford W., Jr. (1976) A Transactions Cost Approach to the Theory of Financial Intermediation. Journal of Finance 31:215-231.

Bommier, Antoine (2006) Uncertain Lifetime and Intertemporal Choice: Risk Aversion as a Rationale for Time Discounting. International Economic Review 47:1223–12246.

Bommier, Antoine (2010) Portfolio Choice under Uncertain Lifetime. Journal of Public Economic Theory 12:57–73.

Bommier, Antoine (2013) Life-Cycle Preferences Revisited. Journal of the European Economic Association 11:1290–1319.

Bommier, Antoine, and Rochet, Jean-Charles (2006) Risk Aversion and Planning Horizons. Journal of the European Economic Association 4:708-734.

Bundorf, M. Kate; Levin, Jonathan, and Mahoney, Neale (2012) Pricing and Welfare in Health Plan Choice. American Economic Review 102(7):3214-3248.

Camerer, Colin (2003) Behavioral Game Theory: Experiments in Strategic Interaction. Princeton, NJ: Princeton University Press.

Chiappori, Pierre-André; Jullien, Bruno; Salanié, Bernard and Salanié, François (2006) Asymmetric Information in Insurance: General Testable Implications. RAND Journal of Economics 37(4):783-798.

Cohen, Alma, and Einav, Liran (2007) Estimating Risk Preferences from Deductible Choices. American Economic Review 97(3):745-788.

Coller, Maribeth; Harrison, Glenn W., and McInnes, Melayne M. (2002) Evaluating the Tobacco Settlement: Are the Damages Awards Too Much or Not Enough? American Journal of Public Health 92(6):984-989.

Cox, James C., and Sadiraj, Vjollca (2006) Small- and Large-Stakes Risk Aversion: Implications of Concavity Calibration for Decision Theory. Games and Economic Behavior 56:45-60.

Di Girolamo, Amalia; Harrison, Glenn W.; Lau, Morten I., and Swarthout, J. Todd (2015) Subjective Belief Distributions and the Characterization of Economic Literacy. Journal of Behavioral and Experimental Economics 59:1-12.

Ehrlich, Isaac, and Becker, Gary S. (1972) Market Insurance, Self-Insurance, and Self-Protection. Journal of Political Economy 80:623-648.

Einav, Liran; Finkelstein, Amy, and Cullen, Mark R. (2010a) Estimating Welfare in Insurance Markets Using Variation in Prices. Quarterly Journal of Economics 125(3):877-921.

Einav, Liran; Finkelstein, Amy, and Schrimpf, Paul (2010b) Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market. Econometrica 78(3):1031-1092.

Elder, Todd E. (2013) The Predictive Validity of Subjective Mortality Expectations: Evidence from the Health and Retirement Study. Demography 50(2):569-589.

Ellsberg, Daniel (1961) Risk, Ambiguity, and the Savage Axioms. Quarterly Journal of Economics 75:643-669.

Epstein, Larry G., and Zin, Stanley E. (1989) Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. Econometrica 57(4):937-969.

Ericson, Keith Marzilli, and Sydnor, Justin (2017) The Questionable Value of Having a Choice of Levels of Health Insurance Coverage. Journal of Economic Perspectives 31(4):51-72.

Fang, Hanming; Keane, Michael P., and Silverman, D. (2006) Sources of Advantageous Selection: Evidence from the Medigap Insurance Market. NBER Working Paper No. 12289, National Bureau of Economic Research.

Fang, Hanming; Keane, Michael P., and Silverman, D. (2008) Sources of Advantageous Selection: Evidence from the Medigap Insurance Market. Journal of Political Economy 116:303–350.

Fang, Hanming, and Wu, Zenan (2018) Multidimensional Private Information, Market Structure, and Insurance Markets. RAND Journal of Economics 49(3):751–787.

Gilboa, Itzhak, and Schmeidler, David (1989) Maxmin Expected Utility with a Non-unique Prior. Journal of Mathematical Economics 18:141-153.

Grether, David M., and Plott, Charles R. (1979) Economic Theory of Choice and the Preference Reversal Phenomenon. American Economic Review 69:623-648.

Handel, Benjamin R. (2013) Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts. American Economic Review 103(7):2643-2682.

Handel, Benjamin R., and Kolstad, Jonathan T. (2015) Health Insurance for 'Humans': Information Frictions, Plan Choice, and Consumer Welfare. American Economic Review 105(8):2449-2500.

Handel, Benjamin R.; Kolstad, Jonathan T., and Spinnewijn, Johannes (2019) Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets. Review of Economics and Statistics 101(2):326-340.

Hansen, Jan. V.; Jacobsen, Rasmus H., and Lau, Morten I. (2016) Willingness to Pay for Insurance in Denmark. Journal of Risk and Insurance 83(1):49-76.

Hansson, Bengt (1988) Risk Aversion as a Problem of Conjoint Measurement. In P. Gardenfors and N-E. Sahlin (eds.) Decisions, Probability, and Utility. New York: Cambridge University Press.

Harris, Katherine and Keane, Michael Keane (1999) A Model of Health Plan Choice: Inferring Preferences and Perceptions from a Combination of Revealed Preference and Attitudinal Data. Journal of Econometrics 89:131-157.

Harrison, Glenn W. (2019) The Behavioral Welfare Economics of Insurance. Geneva Risk & Insurance Review 44(2):137–175.

Harrison, Glenn W; Lau, Morten; Ross, Don, and Swarthout, J. Todd (2017) Small Stakes Risk Aversion in the Laboratory: A Reconsideration. Economics Letters 160:24-28.

Harrison, Glenn W., and List, John A. (2004) Field Experiments. Journal of Economic Literature 42(4):1013-1059.

Harrison, Glenn W.; Martínez-Correa, Jimmy, and Swarthout, J. Todd (2014) Eliciting Subjective Probabilities with Binary Lotteries. Journal of Economic Behavior and Organization 101:128-140.

Harrison, Glenn W; Martínez-Correa, Jimmy, and Swarthout, J. Todd (2015a) Reduction of Compound Lotteries with Objective Probabilities: Theory and Evidence. Journal of Economic Behavior & Organization 119:32-55.

Harrison, Glenn W.; Martínez-Correa, Jimmy; Swarthout, J. Todd, and Ulm, Eric R. (2015b), Eliciting Subjective Probability Distributions with Binary Lotteries. Economics Letters 127:68-71.

Harrison, Glenn W.; Martínez-Correa, Jimmy; Swarthout, J. Todd, and Ulm, Eric (2017) Scoring Rules for Subjective Probability Distributions. Journal of Economic Behavior & Organization 134:430-448.

Harrison, Glenn W.; Monroe, Brian, and Ulm, Eric (2022) Recovering Subjective Probability Distributions: A Bayesian Approach. CEAR Working Paper 2022-03, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University.

Harrison, Glenn W.; Morsink, Karlijn, and Schneider, Mark (2022) Literacy and the Quality of Index Insurance Decisions. Geneva Risk & Insurance Review 47:66-97.

Harrison, Glenn W., and Ng, Jia Min (2018) Welfare Effects of Insurance Contract Non-Performance. Geneva Risk and Insurance Review 43(1):39-76.

Harrison, Glenn W., and Ng, Jia Min (2019) Behavioral Insurance and Economic Theory: A Literature Review. Risk Management & Insurance Review 22:133-182.

Harrison, Glenn W., and Ross, Don A. (2018) Varieties of Paternalism and the Heterogeneity of Utility Structures. Journal of Economic Methodology 25(1):42-67.

Harrison, Glenn W., and Ross, Don (2023) Behavioral Welfare Economics and the Quantitative Intentional Stance. In G.W. Harrison and D. Ross (eds.), Models of Risk Preferences: Descriptive and Normative Challenges. Bingley, UK: Emerald, Research in Experimental Economics.

Klibanoff, Peter; Marinacci, Massimo, and Mukerji, Sujoy (2005) A Smooth Model of Decision Making under Ambiguity. Econometrica 73(6):1849-1892.

Laury, Susan K.; McInnes, Melayne M. and Swarthout, J.Todd. (2012) Avoiding the Curves: Direct Elicitation of Time Preferences. Journal of Risk & Uncertainty 44:181-217.

Loewenstein, George; Friedman, Joelle Y.; McGill, Barbara; Ahmad, Sarah; Linck, Suzanne; Sinkula, Stacey; Beshears, John; Choi, James J.; Kolstad, Jonathan; Laibson, David; Madrian, Brigitte C.; List, John A., and Volpp, Kevin G. (2013) Consumers' Misunderstanding of Health Insurance. Journal of Health Economics 32(5):850-862.

Machina, Mark J. (1987) Choice under Uncertainty: Problems Solved and Unsolved. Journal of Economic Perspectives 1(1):121-154.

Mayers, David, and Smith, Clifford W., Jr. (1981) Contractual Provisions, Organizational Structure, and Conflict Control in Insurance Markets. Journal of Business 54:407-434.

Mayers, David, and Smith, Clifford W., Jr. (1983) The Interdependence of Individual Portfolio Decisions and the Demand for Insurance. Journal of Political Economy 91(2):304-311.

Nau, Robert F. (2006) Uncertainty Aversion with Second-order Utilities and Probabilities. Management Science 52:136-156.

Neilson, William S. (2010) A Simplified Axiomatic Approach to Ambiguity Aversion. Journal of Risk and Uncertainty 41:113-124.

Quiggin, John (1982) A Theory of Anticipated Utility. Journal of Economic Behavior & Organization 3(4):323-343.

Rabin, Matthew (2000) Risk Aversion and Expected Utility Theory: A Calibration Theorem. Econometrica 68:1281-1292.

Richard, Scott F. (1975) Multivariate Risk Aversion, Utility Independence and Separable Utility Functions. Management Science 22(1):12–21.

Richter, Andreas; Schiller, Jörg and Schlesinger, Harris (2014) Behavioral Insurance: Theory and Experiments. Journal of Risk & Uncertainty 48:85-96.

Savage, Leonard J. (1972) The Foundations of Statistics. New York: Dover, Second Edition.

Schmeidler, David (1989) Subjective Probability and Expected Utility without Additivity. Econometrica 57:571-587.

Segal, Uzi (1990) Two-Stage Lotteries Without the Reduction Axiom. Econometrica 58(2):349-377.

Segal, Uzi (1992) The Independence Axiom Versus the Reduction Axiom: Must We Have Both? In W. Edwards (ed.), Utility Theories: Measurements and Applications. Boston: Kluwer Academic Publishers.

Spinnewijn, Johannes (2017) Heterogeneity, Demand for Insurance, and Adverse Selection. American Economic Journal: Economic Policy 9(1):308-343.

Starmer, Chris (2000) Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice Under Risk. Journal of Economic Literature 38:332-382.